Intel Developer Update is Intel's monthly online news magazine for developers. As the official publication of developer.intel.com, it brings hardware, software, and Web developers the latest information on Intel initiatives, technologies, platforms, and products.

Cover Story
Each month, we run a cover story on the most significant industry announcement, trend, or development for the month.

Featured Articles
Delivering in-depth reports on key platforms, products and technologies, our featured articles provide a monthly source of information on issues affecting developers. Be sure to check in every month for the latest developments driving the evolution of the industry.

Contact the Editor
To make *Intel Developer Update* a better information resource, we invite you to share your thoughts on what we've published or what you'd like to see covered. Comments are always welcome.

Archives
Our archives contain two groups of previously published articles. One group contains all the articles that appeared in *Platform Solutions News*, the earlier version of *Intel Developer Update*. The articles date from September 1997 through August 1999. The other group is set up to contain *Intel Developer Update* articles dating from the inaugural September/October 1999 issue.

Bookmarking
We advise against bookmarking article pages. They're accessible online only during the month the issue is live. Thereafter, they're removed to our archives. Instead, we suggest that you bookmark the PDF (Adobe* Portable Document Format) file versions of the articles. You'll find buttons for the PDF files labeled "print article" in the right navigation section of each article. A PDF for the entire issue is labeled "print magazine" and is located near top right side of the IDU home page.

On behalf of all of us at Intel Developer Update, welcome to the future of the PC platform!

Table of Contents

(Click on page number to jump to articles)

Note: Intel does not control the content on other company's Web sites or endorse other companies supplying products or services. Any links that take you off of Intel's Web site are provided for your convenience.

Cover Story

Transitioning to Gigabit Ethernet

Blaine Kohl
Director of Product Marketing
LAN Access Operation
Intel Corporation

Overview

To meet the growing demand for additional bandwidth in new applications, high-speed connections are required to reduce traffic bottlenecks and improve performance on network systems. At the same time, higher bandwidth solutions must be compatible with existing technologies to protect large investments in network infrastructure.

According to In-Stat, a national market research firm, by the end of 2000, 83.3 percent of all Ethernet NICs (network interface cards) shipped will be 10/100 Mbps. Two percent will be 1000 Mbps and the remaining 10.2 percent will be 10 Mbps. By 2004, it is expected that Gigabit Ethernet will be 59.2 percent of all Ethernet NICs shipped (In-Stat, Networking LAN Service, report number LN9914NC).

In addition to end users' need for more bandwidth, more and more users are now being networked together. According to In-Stat, in 1997 some 41.8 million NICs shipped worldwide, and that number jumped to 52.7 million in 1999 (In-Stat, Networking LAN Service, report number LN9914NC). The growth of Gigabit is also apparent in the compound annual growth rate (CAGR). Gigabit Ethernet is expected to have a 49.3 percent growth through the year 2003 (International Data Corp, Worldwide LAN Semiconductor Market Review and Forecast, 1997–2003, report number: W20947–December 1999).

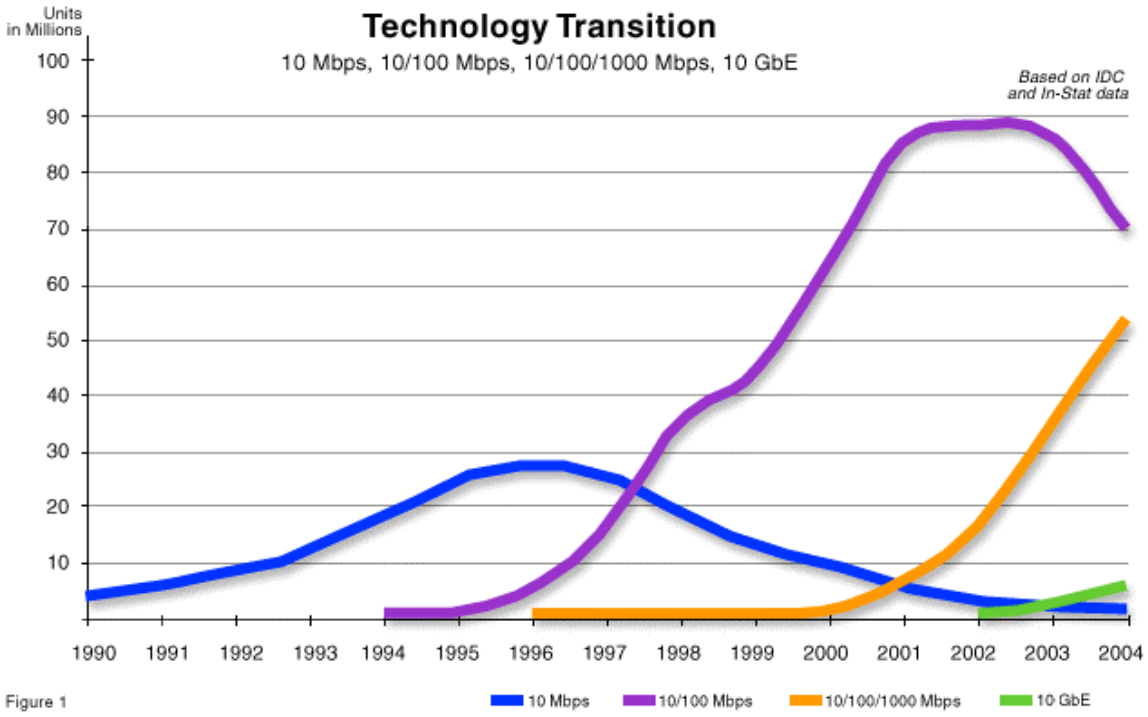

Figure 1

These two trends created an exponential increase in network bandwidth usage. For example, today a typical 100-node network may require 70 users connected at 10 Mbps and 30 at 100 Mbps, making the total bandwidth 3,700 Mbps. If 20 more stations are added over the next 12 months, of which 15 need 10 Mbps and 5 need 100 Mbps, the network must supply 650 Mbps additional bandwidth for a total of 4,350 Mbps, or an 18 percent increase in needed bandwidth.

To meet these increasing requirements, Gigabit Ethernet has become the industry-standard solution for high-speed local area networking. In addition to meeting the additional bandwidth requirements, Gigabit Ethernet (1000BASE-T) offers a cost-effective solution with both scalability and protection of investment.

Gigabit Ethernet runs traditional Ethernet frames and extends current 10-Mbps and 100-Mbps Ethernet networks seamlessly. This combination lets network administrators upgrade networks to 1000 Mbps data rates while maintaining the simplicity and manageability of Ethernet networks, just as 100-Mbps Ethernet extended 10-Mbps Ethernet networks.

Getting to 10/100/1000 Mbps

Proponents of Fast Ethernet quickly learned that adoption would proliferate best with a migration strategy offering both Ethernet and Fast Ethernet speeds. The same has been applied to Gigabit Ethernet by offering 10/100/1000-Mbps solutions.

One of the greatest strengths of 10/100/1000-Mbps Ethernet is the ability to integrate seamlessly into current Ethernet network topology, where immediate speed increases are needed. The migration from 10/100-Mbps Ethernet to 10/100/1000-Mbps Ethernet is important because it allows users an immediate upgrade to currently congested areas.

The server link and the backbone are the two network areas under the greatest strain in current networking environments. Gigabit Ethernet fits well into both these categories. By offering the ability to upgrade to higher bandwidth in the future, 10/100/1000-Mbps Ethernet also allows users headroom with links currently operating at 100 Mbps. Workstations and high-performance desktops are a prime example of this need. Current workstations are linked with 100-Mbps Ethernet. That link may be sufficient now, but workstations will exceed the capacity of 100-Mbps Ethernet in the near future.

To meet growing needs for additional bandwidth, the IEEE P802.3z standard specifies 1000-Mbps-Ethernet operation over fiber optic cable, while the IEEE P802.3ab standard enables 1000-Mbps operation over Category 5 (CAT-5) or higher rated copper cable. This allows 1000-Mbps Ethernet to operate over the most commonly installed cabling.

With the variety of copper and fiber choices available today, network managers are asking whether legacy Category 5 cabling will support 1000BASE-T Gigabit Ethernet. According to industry experts, any links currently using 100BASE-TX should easily support 1000BASE-T without replacing Category 5 cabling.

Gigabit Cu Implementation

100BASE-TX achieves 100 Mbps by using 4B5B coding to send three-level binary encoded symbols across the link at 125 Mbaud. Gigabit Ethernet over copper (1000BASE-T) uses four pairs of two twisted cables (one to transmit and one to receive) to transmit at a 125-Mbaud rate with a five-level coding scheme for each link. Because 1000BASE-T sends and receives simultaneously on each pair, the combination of five-level coding and four pairs allows 1000BASE-T to send one byte in parallel at each signal pulse.

With Category 5 cabling installed according to the specifications in ANSI/TIA/EIA/568A, it is unlikely that more than 10 percent of existing cable was installed improperly. Since improperly installed cable would not have supported 100BASE-TX, any installation-related problems were most likely discovered and corrected during deployment. However, it is important that network managers test cabling for the effects of insertion loss and link-induced interference in 1000BASE-T caused by echo and crosstalk. Installed cabling can be field-tested by using automated return loss and ELFEXT tests incorporated into current versions of cable test tools.

Cost-effective Cu Solutions

Unlike FDDI (Fiber Distributed Data Interface) and ATM (Asynchronous Transfer Mode), Gigabit Ethernet addresses the bandwidth dilemma without requiring costly infrastructure changes. Given the difficulty in replacing horizontal cabling, network managers make cable decisions with expectations that the cable solution will offer the flexibility for increasingly complex applications.

A 10/100/1000BASE-T solution offers:

- Increased network performance levels, including traffic localization and high-speed cross segment movement that reduce bottlenecks.
- Increased network scalability to manage more users and bandwidth-hungry applications.
- Decreased costs overall, as copper is less expensive and easier to maintain than fiber.

Deployment Scenarios

Because 1000BASE-T solutions offer compatibility, they offer IT managers a more seamless migration path to eliminating backbone bottlenecks in desktops today. This solution will be available in workstations by 2002 and desktop deployment by 2003. These steps are likely to occur along the way:

Gigabit Ethernet will be switched and routed at the network backbone with switch-to-switch connections beginning in mid-2000, due to a high-volume ramp in 1000BASE-T delivery.

Switch-to-server deployments will be implemented next, boosting access to critical server resources. In addition, 10/100/1000-Mbps implements allow operation at 100 Mbps until that portion of the desktop is ready to migrate to 1000 Mbps.

Finally, as desktop connectivity costs decrease, as Gigabit Ethernet has proliferated throughout the backbone, and network demands continuously increase, Gigabit Ethernet will enter the desktop space in 2003.

Summary

As an extension of current Ethernet technology, Gigabit Ethernet is compatible with the most widely deployed networking architecture. In particular, copper Gigabit Ethernet implementation compatibility preserves investment in system administration, requires minimal training, involves little or no need for additional protocol stacks, and requires no investment in new cabling. Most important, Gigabit Ethernet can efficiently match the power of high-performance networks, providing the needed infrastructure from the desktop to the backbone.

The performance and speed advantages of Gigabit Ethernet technology, combined with the ease of migration through use of existing network infrastructure, offer compelling motivation for adopting 1000BASE-T solutions.

More Info

For more information on Gigabit Ethernet products, visit the Intel Networking and Communications Web site.

For information about an open industry forum whose purpose is to promote industry cooperation in the development of Gigabit Ethernet, visit the Gigabit Ethernet Alliance Web site . The site features a white paper section that takes a look at the new Gigabit Over Copper, 1000Base-T cabling, and other related topics.

Author Bio

Blaine Kohl is the director of product marketing for Gigabit Ethernet Products in Intel's LAN Access Operations. Blaine has held positions at Intel, Level One Communications, Jato Technologies, and Motorola, Inc.

Columns

Inside Looking In

Grounded in Good Management

Tim Mostad
Senior Technical Marketing Manager
Intel Corporation

Column

I recently won an Intel divisional award for outstanding people management. While I'm happy about it, I don't consider management an extraordinary achievement. To me management is simply the process of unleashing the potential of the people who do the real work. Management done well is invisible and, conversely, management done poorly is obvious. In fact, spending time drawing attention to management serves to undermine the effectiveness of the overall organization. If the organization is getting things done, why ruin its effectiveness by calling attention to the person at the helm?

The most obvious answer is that we need role models for others to emulate. Good concept, but as anyone who's survived the standard Intel management training regimen knows, management is situational. What makes a manager good in one situation would make him or her completely ineffective in another. Emulation is an otherwise good mechanism that, when simply applied, will lead to inconsistent results.

A newer, more advanced Intel management-training course introduces the idea of versatility and the need to be aware of not only the situation but also your style so you can adjust it to be more effective. Knowing that the *Titanic* was sinking, the captain had to behave differently than he would refereeing the shuffleboard tournament. He focused on what needed to be done and undoubtedly saved some lives, but he ended up going down with the ship anyway. The lesson here is that as a manager you can be in a situation where, regardless of what you accomplish, there will be no way to save yourself. However, without versatility a manager can sink a ship even in calm water on a clear day. So the second reason for management awards might be as a commendation for an often complex and sometimes hopeless job.

I think there's a third and more important reason: to let the real workers know that the company pays attention to the way the managers ply their trade. Knowing that higher-level forces with the power to affect your paycheck are focusing on your work is a pretty effective behavior modification mechanism for most managers. And a ray of hope for the folks in their chain of command.

If nothing else, maybe this type of award will raise the "sandbar" just a little bit, high enough to rid companies of the "nightmare manager." You know the person; every organization of any size has at least one. A lot of individual characteristics can contribute to this label. At Intel, one particular type seems to fit nicely into a niche of our system, and I've yet to see a remedy.

They're among the take-charge go-getters. Intel culture promotes this style, and why not? Everyone wants to work with someone who gets things done. However, for some of these managers human collateral damage is ignored as they blunder through the organization like a battleship without a rudder.

A telltale sign you're dealing with one of these managers is reliance on the consensus process whenever significant risk presents itself. Why? If something goes wrong, blame can efficiently be deflected. This behavior is fully rationalized in the name of communication. After all, nobody should ever make a decision without involving those affected, right? Our prototypical bad manager turns his guns toward his own people, the ones who should have warned him way from the action that failed.

In fact, faultfinding and blame attribution are part of the bad manager's arsenal. When these managers encounter anything they're not associated with, they typically use "analysis paralysis" to filibuster it long enough to find a problem, which they in turn "fix" and use to assume control. They like to characterize themselves as troubleshooters. The term couldn't be more appropriate because they can smell trouble like blood in the water and create trouble where there is none. Then they call great attention to themselves as they shoot the trouble dead. The quiet influence that makes others successful is not their style.

Instead of rewarding good managers, maybe we should engage in public humiliation of bad managers. I know it sounds cruel and unusual, but how else do you deliver a strong enough message? We may not be able to afford the time to let the power of positive recognition slowly take effect. Maybe a badge of disgrace stitched on bad managers' clothes is the answer? We could pick a really loud color and bring them all the recognition they seem to crave, if not exactly the kind they want.

In the Navy, a ship captain's career effectively ends by simply grounding a ship on a sandbar, even when there's no damage to the ship and it's just stuck in soft sand. Since sandbars shift constantly with the tides and are often hidden by shallow water, careers are in constant jeopardy. This punishment sends a very strong message to managers who have responsibility for hundred-million-dollar machines and the lives of hundreds or thousands of sailors. Something more moderate needs to be in place to let high-tech organizations benefit from the learning that happens day-to-day.

So maybe positive recognition really is the best answer. The organizational win happens when a lot of good people accomplish great things, not the least of which may be recognizing that they're all sailing forward together, and more or less on course.

Author Bio

Tim Mostad continues to pursue technical marketing nirvana by applying his 19 years of Intel hardware experience to extending Intel's influence with software and Internet developers. As operations manager in Intel's Developer Relations Division, Tim focuses on the development of broad and efficient enabling processes and infrastructure, primarily through use of the Internet.

From the Editor

Donna Loveland
Managing Editor
Platform Marketing
Intel Corporation

Column
How do you manage?

That question appears to be replacing "How are you doing?" as the greeting of choice for this decade. At least I'm hearing it a lot, not just in cubicle doorways but also in the aisles of trade shows and the lobbies of office buildings and the other places where we live our working lives.

It's so like us tech-types to come up with an expression that efficiently combines salutation, sympathy, and search for information all at once. The reason it's caught on, though, is that it reflects the state we're in relative to technology, just as "going nonlinear" seemed so apt a few short paradigm shifts ago.

All of us, individuals and corporations alike, are taking on more than ever. We have to. Our work is now as diverse as the rest of our culture. No one can prosper in a discipline like communications and focus on chip-level issues to the exclusion of industry developments in broadband, wireless, Internet, and a score of other technologies with subtechnologies of their own.

So how does Intel manage?

The latest Intel Developer Forum conference offered a pretty strong answer. To take on more, you take on more.

Yes, Intel will continue doing silicon and doing it superbly … along with e-Business and e-Homes and e-gad, maybe even an order of fries with that. The philosopher Teillhard de Chardin, who's also gaining currency, once observed that everything that rises must converge—that those disparate elements that come up in daily life and later become top-of-mind do eventually come together in a sensible, integrated whole. You manage by keeping your eyes on the prize at the pinnacle of that convergence.

In this issue, "Inside Looking In" columnist Tim Mostad examines what it takes to lead in a demanding environment in **"Grounded in Good Management."**

With **"Intel Drives In-Vehicle Solutions,"** Lee Machen offers a look at how audio, video, wireless communications, and mundane transportation are merging in highly versatile personal vehicles planned to arrive in consumers' driveways soon.

"Scalable Board Design= Lower Total Cost" focuses on a program whose reason for being is consolidation. The Intel Scalable Performance Board Design Program provides a solution for a wide variety of applied computing applications by supporting different processor and chipset technologies in a single board design.

Wide-ranging applications from industrial laptops to Formula One race cars to ATMs to Internet routers need to operate reliably whatever the temperature constraint. **"Thermal Management for Applied Computing"** describes how a well-designed hardware thermal solution combined with software thermal management preserves system functionality and extends device life cycle.

Looking forward, the newly formed Serial ATA Working Group seeks to provide the framework to enable future storage technologies by providing a new ten-year storage interface roadmap. Steve Spina's **"New Serial ATA Working Group Formed"** has the details.

Last but not least, in our cover story, **"Transitioning to Gigabit Ethernet,"** Blaine Kohl takes a focused look at bringing faster connections and higher performance to networked systems without the costly retrofit of legacy cabling. If you want to understand the performance and speed advantages of Gigabit Ethernet technology, this article is for you.

We hope you'll manage to find material well worth your while in the array published here in April's *Intel Developer Update*.

Author Bio

Donna Loveland is the editor of *Intel Developer Update* magazine. She joined Intel's Platform Marketing group in 1999 as the editor of Platform Solutions News. Donna began her career with Intel in 1982 as a technical editor in an advanced microprocessor development group. Since then, she's held technical and marketing positions in leading-edge technology areas ranging from stereoscopic display to digital broadcast to scalable online content. Donna has a B.A. degree in English from the University of Rochester and an M.A. in Expository Writing from the University of Iowa.

Departments

Applied Computing

Scalable Board Design=Lower Total Cost

Tal Howell
Technical Marketing Engineer
Embedded Intel Architecture
Intel Corporation

Greg Brown
Technical Marketing Engineers
Embedded Intel Architecture
Intel Corporation

Bill Wilson
Product Marketing Engineer
Embedded Intel Architecture
Intel Corporation

Overview

What if you could use one scalable board design for different processors to meet the needs of a wide range of embedded applications?

The Intel® Scalable Performance Board Design Program for applied computing provides the advantage of a single board design with either the Intel® 440BX chipset or Intel® 810 chipset. Boards using these designs can be populated at build time with Intel® Celeron™ or Intel® Pentium® III processors without board changes. These board designs meet a variety of customer requirements by supporting application-specific software and peripherals, while maintaining compatibility with operating systems, BIOS, and peripheral interfaces. The Intel Scalable Performance Board Design Program also:

- Provides price and performance options
- Reduces the design and validation effort for multiple designs resulting in faster time-to-market
- Lowers total cost of ownership by reducing warehouse inventory and manufacturing costs, and debug and technical support costs

Price and Performance Options

The Intel Scalable Performance Board Design Program currently supports Celeron processors in the PPGA (Plastic Pin Grid Array) package and Pentium III processors in the FC-PGA (Flip-Chip Pin Grid Array) package, providing a wide variety of price and performance options in a single board design.

The Celeron processor (PPGA) features an integrated 128-Kbyte level-two cache with a separate 16-Kbyte instruction and 16-Kbyte data level-one cache and is available in a range of speeds from 300 to 433 MHz with a 66-MHz bus.

The Pentium III processor (FC-PGA) features an integrated on-die, 256-Kbyte, 8-way set associative level-two (L2) cache. The L2 cache implements the new Advanced Transfer Cache Architecture with a 256-bit-wide bus. This processor also includes a 16-Kbyte level-one (L1) instruction cache and 16-Kbyte L1 data cache. These cache arrays run at the full speed of the processor core. The level-two cache is capable of caching 64 Gbyte of system memory address space. The Pentium III processor in the FC-PGA package is currently available at 600 MHz with a 100 MHz bus.

The Intel Pentium III processor for the PGA370 socket uses the same core and offers the same performance as the Intel Pentium III processor for the SC242 connector. This package fits the same 370-pin zero insertion-force socket (PGA370) used by the Celeron processor. The heat sink is attached directly to the back of the processor core package without the use of a thermal plate or heat spreader.

440BX Chipset Board Design

The Intel 440BX Chipset has been validated for use with both Celeron and Pentium III processors and provides a high level of performance while containing total system cost. The 440BX chipset incorporates ECC memory control to provide a high level of data integrity, critical for maintaining uptime in applications.

The physical interface design for the Intel® 82443BX PCI/AGP controller is based on the GTL+ (Gunning Transceiver Logic) specification. It requires open-drain buffers and pull-up resistors to terminate the transmission line and control the rise and fall times of signals. AGTL+ (*Assisted* GTL+) integrates the termination resistor in the processor.

The Intel 440BX chipset includes:

- Host-to-PCI bridge
- Optimized DRAM controller and data path
- Accelerated Graphics Port (AGP) interface based on the AGP Specification Rev 1.0 to support up to 133-MHz data transfer rates for 3D graphics and video applications
- Interface to the PIIX4E, the I/O subsystem portion of the PCI/AGP platform
- Optimized 72-bit DRAM interface (64-bit data plus ECC) that supports 3.3V DRAM technology
- Interface to a PCI bus operating at 33 MHz (compliant with PCI Rev 2.1 Specification) providing a high-performance component-level interconnect

810 Chipset Board Design

The Intel 810 chipset is a scalable solution supporting both a 66-MHz and 100-MHz front side bus, and a 100-MHz memory interface. It has been validated for use with both Celeron and Pentium III processors. The graphics accelerator architecture consists of dedicated multimedia engines executing in parallel to deliver high-performance 3D, 2D, and motion-compensation video capabilities. An integrated centralized memory arbiter allocates memory bandwidth to multiple system agents to optimize system memory utilization. A new chipset components interconnect, the Accelerated Hub Architecture, is designed into the Intel 810 chipset to provide an efficient communication channel between the Graphics Memory Controller Hub and the I/O controller hub.

The Intel 810 chipset includes:

- A new security and manageability infrastructure through the Firmware Hub component
- Support for the Full-on (S0), Stop Grant (S1), Suspend to RAM (S3), Suspend to Disk (S4), and Soft-off (S5) power-management states
- Supports Alert-on-LAN* for remote administration and troubleshooting
- Elimination of ISA provides true plug-and-play
- Addition of AC ’97 allows the OEM to use software-configurable AC ’97 audio and modem coder/decoders (codecs) instead of the traditional ISA devices

The Intel 810 chipset contains three core components: the Graphics Memory Controller Hub (GMCH), the I/O controller hub, and the Accelerated Firmware Hub (FWH). In the host controller, the 82810 Graphics and Memory Controller Hub integrates a 66/100-MHz system bus, 100-MHz SDRAM controller, and a high-speed Accelerated Hub Architecture for communication with the I/O Controller Hub. The Intel 810 chipset also features Digital Video Out Port for LLD and TV out support, and Intel® DVMT Graphics Architecture. The Intel 810 chipset saves overall system cost by integrating graphics into the GMCH.

Board Design Considerations

The design guides for the 440BX and 810 chipsets provide the information and design recommendations required to ensure compatibility with either Celeron or Pentium III processors. Areas of focus include AGTL+ signals and termination, V_{REF} inputs, system bus clock, CMOS compatibility, processor core voltage decoupling, Phase Lock Loop (PLL) power considerations, bus frequency selection, and definitions of new pins. A large section of the design guide is devoted to important board layout and routing guidelines, including AGTL+ layout recommendations, pre- and post-layout simulations, timing analysis, and host clock routing spacing. Board design guidelines are also provided for DC, system bus AC, and thermal considerations, in addition to PGA370 socket guidelines. The design guides also include a design checklist and third-party vendor information.

Summary

The Intel Scalable Performance Board Design Program provides a versatile solution for a wide variety of applied computing applications based on Embedded Intel® Architecture. By supporting different processor and chipset technologies in a single board design, designers receive price and performance options and reduced design and validation effort for multiple designs. This results in faster time-to-market, lower total cost of ownership by reduced warehouse inventory, manufacturing costs, minimal debug, and technical support costs.

More Info

For technical information, refer to:

- Intel® 440BX AGPset/PGA370 Scalable Performance Board Design Guide
- Intel® 810 Chipset Scalable Performance Board Design Guide
- Pentium® III Processor for the PGA370 Socket up to 750 MHz Datasheet
- P6 Family of Processors–Hardware Developer's Manual

Author Bio

Tal Howell is a senior technical marketing engineer at Intel. He provides technical support for the Transaction Terminal Platform. During his 19 years with Intel Tal has been involved with Intel Architecture training. For three years he served as curriculum manager for To-the-Point Technology Training, which focused on IA-32 embedded designs. Tal attended DeVry Institute of Technology.

Greg Brown is a technical marketing engineer at Intel. He is responsible for providing technical support for the 810 and 815 chipsets for applied computing. Greg joined Intel in 1996 and spent the past three years in the Intel Quality and Reliability group as a development engineer. Greg graduated in 1996 from Arizona State University with a bachelor's degree in Electrical Engineering.

Bill Wilson is a product marketing engineer at Intel. He has product management responsibilities for Intel Celeron processors and the Intel 810- and Intel 440BX chipsets. Bill joined Intel Manufacturing in 1984. He holds a B.S.B.A. degree from the University of Phoenix.

Thermal Management for Applied Computing

Anna Madrid
Sr. Technical Marketing Engineer
Embedded Microcomputer
Intel Corporation

Prosenjit Ghosh
Senior Technical Staff
Mobile Computing Group
Intel Corporation

Michel Goulet
Director of Software
Teknor Applicom

Yves Desrochers
Engineering Group Director
Teknor Applicom

Overview

Advances in microprocessor technology have enabled system developers to provide abundant processing power in increasingly small, highly integrated devices. A benefit of reduced processor size is reduced power consumption. But with higher performance, thermal management needs to be addressed. In particular, applied computing devices, which often operate in warm or dry environments and closed system cages or chassis, are sensitive to this need. Such environments include industrial laptops, industrial process control systems in a nuclear station, even function regulation systems for a Formula One race car. Point-of-sale, ATM, and kiosk applications also require thermal management. In addition, communication hardware (routers, switches, appliance blades, server applications, etc.) also requires reliable operation under high temperatures or other external constraints. A well designed hardware thermal solution combined with software thermal management preserves a system's functionality and extends the life cycle of an electronic device.

Fundamentals of Thermal Design

Thermal design is deceptively complex. A successful design requires a thorough understanding of thermodynamics, fluid mechanics, and heat transfer. Unlike most mechanical engineering problems, thermal behavior can be subtle and transparent—sometimes counter-intuitive—and difficult and time-consuming to comprehend. Ignoring thermal performance can severely affect component or device reliability, timing circuitry, and system performance.

According to thermodynamic concepts, when energy is transferred as a result of temperature difference, it is called heat. Basic thermodynamic theory helps explain and deal with energy transfer in systems in equilibrium or systems moving from one state of equilibrium to another. However, it does not address the rate at which this energy transfer occurs. The science of heat transfer supplements the first and second laws of classical thermodynamic theory by providing analytical methods to predict rates of energy transfer.

Energy in the form of heat is transferred through one of three modes or a combination of these modes: conduction, convection, and radiation. Conduction heat transfer occurs when one solid medium contacts another medium and a temperature gradient exists between the two bodies.
When a moving fluid comes in contact with a solid surface at a different temperature, the resulting thermal-energy exchange process is called convection heat transfer. There are two kinds of convection processes: natural, or free, convection and forced convection. Free convection occurs when the motive force comes from the density difference in the fluid, which results from its contact with a surface at a different temperature and gives rise to buoyant forces. Forced convection occurs when an outside motive force moves a fluid past a surface at a higher or lower temperature than the fluid.

Whereas conduction and convection heat transfer can take place only through a material medium, radiation can transport heat even through a perfect vacuum. In the radiative mode of heat transfer, the energy is transported in the form of electromagnetic waves that travel at the speed of light.

Thermal Management Solutions

Understanding the thermal behavior of a system and optimizing that behavior for cost efficiency requires modeling and analyzing the system. Component reliability is inversely proportional to temperature gradients and sustained exposure at the elevated temperatures. Thus, power dissipation is a key function of any thermal solution.

Effective thermal management requires identifying critical issues such as shock, vibration, and usage scenarios, considering potential coupling effects, and developing an integrated, interdisciplinary solution. Successful applied computing solutions today utilize a combined hardware and software approach. System hardware solutions are based on internally distributing and externally dissipating the thermal energy to the ambient environment. System software solutions typically regulate the dissipative power of the device based on active system feedback controls.

Hardware

Developing a successful thermal solution requires hardware. Such hardware consists of the use of heat sinks, fans, and heat pipes as well as interface materials, heat sink mounting, allocated system airflow mechanisms, fans, chassis ventilation, and component and motherboard placement.

A heat sink cools a device by expanding the surface area of the part to which it is attached, increasing the amount of heat that can be cooled by the ambient air. In evaluating heat sinks, consider the following: cost, extruded versus corrugated (folded) fin, extruded versus bonded fin, flexibility in x, y, and z dimensions, and system airflow.

Fan solutions include passive-active and active approaches. With passive-active, fan heat sink solutions provide airflow and require little or no system airflow, while active fan heat sink solutions incorporate a fan that is attached to the solution. The airflow that a fan produces blows parallel to the fan's blade axis. In considering these solutions, points to evaluate are similar to those for heat sinks: cost, flexibility in x, y, or z dimensions, system airflow, and reliability.

Heat pipes, in their simplest sense, are heat movers or spreaders. They acquire heat from a source and move or spread it to a region where it can be more readily rejected. Considerations in using heat pipes include: single-slot CPCI z-height, power consumption, noise elimination/reduction, low maintenance, sealed enclosure cooling, and extended ambient temperatures.

Software

As processor speed and power consumption have increased, so has the need for temperature control. Originally, temperature control on a board was achieved by using the system BIOS and "passive cooling" functionality. Passive cooling reduces the speed or disables some on-board devices in order to decrease power consumption, thereby reducing overall system temperature. "Active cooling," on the other hand, increases power consumption by activating fans that increase airflow, thereby reducing the system temperature.

In a software-based approach, the system will display CPU die temperature in real time, allowing the user to read the normal temperature of the system and set accurate values for resume and overheat.

With an Advanced Configuration and Power Interface (ACPI) solution, temperature control is moved from the BIOS to the operating system. In keeping with more sophisticated temperature control features, the inclusion of the ACPI in operating systems lets the user make more intelligent decisions with a better follow-up on the CPU load and applications control.

The ACPI thermal design is based on regions called thermal zones. Some systems may have more than one thermal sensor to allow subdividing the system in many more thermal zones.
A benefit of ACPI is that it standardizes thermal control methods. In contrast, BIOS control methods are proprietary and do not let applications use them transparently. ACPI control also provides the choice of passive cooling, active cooling, or a mix of both.

Summary
Depending on the system environment, available power, and available airflow, temperature control can be exercised more precisely through reliable thermal management techniques. A combination of hardware and software thermal solutions is key to successfully managing thermal needs. A well-adapted thermal management implementation increases system reliability and creates a more stable environment for applied computing applications.

More Info
You can find numerous sites and documents on the Intel Developer Web site. These areas within the site offer specifics you may find helpful:

- The Applied Computing Solutions Guide
- The Packaging Databook
- Intel Architecture Labs
- Pentium® II Processor–Low Power Module Thermal Design Guide

Author Bios
Anna Madrid is a senior technical marketing engineer with 10 years of experience at Intel. After a year in the company's Graduate Rotation Program Anna joined Intel's Assembly Technology Development Group. Since 1997 she's worked in the Applied Computing and Embedded Microcomputer Division, where her current research focuses on optical, electrical, and thermal technology development, semiconductor package technology, and manufacturing. Anna received a B.S.E.E. degree from the University of Arizona. She received her M.S.E.E. degree from Arizona State University.

Prosenjit Ghosh is a senior technical staff member with Intel's Mobile Computing Group, where he directs thermal technology development for the Mobile Module Operation. Formerly with Texas Instruments' Mobile Computing Architecture Lab, Prosenjit's research interests include mobile and portable electronics system thermal design, thermal interface materials, and nanoscale technologies. He has authored and co-authored numerous papers and presentations. Prosenjit is an alumnus of the Mechanical Engineering programs at New Jersey Institute of Technology. He is an active member of ASME and IEEE.

Michel Goulet is director of software at Teknor Applicom. His team specializes in software project management for low-level BIOS designs and drivers for Teknor single board computers and integrated systems. He has recently published several papers on embedded software design. Before joining Teknor, Michel was professor of mathematics and computer science and assisted in real-time, multi-agent computing architecture research. He holds a Computer Engineering degree from the University of Laval in Quebec City, Canada.

Yves Desrochers is engineering group director at Teknor Applicom and currently leads Teknor's PICMG and VIPer single board computer engineering team. Yves worked at Matrox and Roctest prior to joining Teknor, and he has more than 10 years of software and hardware experience in computer design, video, and instrumentation. He holds an E.E. degree from Montreal's Ecole Polytechnique.

Desktop

New Serial ATA Working Group Formed

Steve Spina
Director of Marketing
Intel Architecture Labs
Intel Corporation

Overview

With an expected initial transfer rate of 1.5 Gbps, scalable to 2X and 4X, Serial ATA is expected to deliver the performance required to accommodate future storage device data rates over the next 10 years. In addition, Serial ATA is expected to anticipate silicon trends with lower pin count and reduced voltage. The efforts of the newly formed Serial ATA Working Group seek to provide the framework to enable future storage technologies by providing a new 10-year storage interface roadmap. The roadmap is planned to link emerging system and storage technologies while maintaining backward compatibility for software.

To facilitate the transition from parallel ATA to Serial ATA, the specification will use the existing ATA infrastructure, allowing current operating systems and application software to run without modification. This approach preserves the investments software developers have made over the last decade.

The Serial ATA Working Group

At the Intel Developer Forum Spring 2000, Intel led the formation of the Serial ATA Working Group, joining with other core promoters including APT, Dell Computer Corp., IBM Corp., Maxtor Corp., Quantum Corp., and Seagate Technology. The mission of the Serial ATA working group is to define the next-generation storage interface that provides scalability for next-generation storage products. The resulting Serial ATA specification will be designed to deliver high bandwidth that cannot be afforded by existing parallel interfaces but will provide backward compatibility by utilizing existing infrastructure.

To meet this objective, the Serial ATA Working Group is aiming to create a Serial ATA Interface specification that is a robust, comprehensive, and extensible industry specification defining the serial interface between storage devices and high-performance PCs. The Serial ATA specification will address protocol and electrical and mechanical definitions. It will be the first specification expressly designed for serial storage devices on PCs and the first to design in a degree of backward compatibility with existing storage technologies.

Developers who implement products compliant with the Serial ATA specification are expected to be able to save costs by eliminating pin counts and cable expenses while laying the groundwork for the 10-year storage interface roadmap.

Royalty-free Licensing

Companies who adopt the Serial ATA specification must agree to royalty-free, reciprocal licensing of essential patents and other forms of intellectual property needed to implement the interface portion of the specification.

Time to Get Involved

Dramatic new storage products are just beginning to emerge. Developers interested in becoming involved should take this opportunity to join Intel and other members of the Serial ATA as contributors. Contributors can deliver input for the development of the Serial ATA specification. By becoming Serial ATA Adopters, they can use the specification to implement the design of Serial ATA-compliant products.

Summary

The Serial ATA interface specification is expected to enable the industry to deliver higher quality storage products in innovative form factors at lower cost. These products will further enhance Ease-of-Use and Visual Computing experiences for users of Intel® Architecture-based PCs. The newly formed Serial ATA Working Group seeks to provide the framework and roadmap for the specification. The Serial ATA Working Group is open to all interested developers who agree to the terms and conditions of membership.

More Info

For more information about the Serial ATA Working Group, the Serial ATA specification, and the benefits of Serial ATA membership, visit theATA Serial Working Group Web site .

Author Bio

Steve Spina is a director of marketing at Intel Corporation, where he focuses on the Scalable Platform Initiative in the Intel Architecture Labs. Steve joined Intel in 1981 and has 12 years of Intel marketing experience in areas ranging from products to technology initiatives. Most recently, Steve served as a strategic initiative manager specializing in the graphics and digital display technologies that support Intel's Ease-of-Use and Visual Computing initiatives. He holds a degree in business management from Linfield College.

Initiatives and Technologies

Intel Drives In-Vehicle Solutions

Lee Machen
Technical Marketing Engineer
Handheld Components Division
Intel Corporation

Overview

"We will do nothing short of transforming our cars and our trucks into portals for the Internet."—Jacques Nasser, CEO, Ford Motor Company, *USA Today*, 2/1/00.

Computing functionality is entering automobiles in ways that will enhance productivity, safety, communications, and entertainment while on the road. Drivers will be able to remain connected to e-mail and the Internet, obtain driving directions, and use their voices to dial a cell phone, all while keeping their hands on the wheel and their eyes on the road.

While the in-vehicle PC seamlessly operates each of these applications and many more, it will also be able to monitor car diagnostics and call for help in the event of an accident. These capabilities are not in the distant future; they are being designed and implemented today. With the wide range of applications being developed for the in-vehicle computing market segment, Intel® products are the best choice to provide the necessary processing power and design flexibility.

Applications

The Big Three auto makers in the US, as well as major companies in Europe and Asia, are evaluating what their in-vehicle computing strategy will be. In the next few years every car will have some level of computing functionality, from basic safety and security in the low-end to more driver- and passenger-oriented functions such as navigation and DVD movies in high-end vehicles.

The most popular application today is navigation using a global positioning system (GPS). This system maps a course to the driver's destination, and gives turn-by-turn directions according to an on-board map database. In the future, maps and updated traffic information downloaded from the Internet through a wireless link might replace these databases. Real-time information could provide an instant update on a road under construction or an accident causing backups on the highway, allowing the system to offer an alternate route.

Entertainment is another important function of the Car PC. Many systems will handle the current car stereo functionality, including radio, CD, and CD changer controls, while adding MP3 player capabilities. Higher-end systems will support backseat passenger entertainment in the form of DVD movie players with surround sound and video games.

Surveys indicate that over 85 percent of cell phone owners use their phones while driving, and integrating them into the car increases both their usefulness and safety. With integrated phones, drivers never need to take their hands off of the wheel because calls can be made through voice commands, and the computer automatically lowers the volume of other audio components for incoming calls.

The cell phone will also be a way for drivers to obtain wireless data in the car, letting them stay connected on the road with e-mail and Internet information. Messages, sports scores, stock quotes, or any other data that is received can then be converted to speech and recited to the driver. Intel's wireless technology roadmap, including processors, baseband chipsets, DSPs, and flash memory, will play an important role.

Safety and security functions will also be implemented in systems across the entire price range of cars. The Strategis Group's November 1999 study indicates that revenues from automotive telematics services are projected to rise from less than $40 million currently to more than $1.7 billion by 2004. Capabilities might include notifying authorities of vehicle location in the case of a collision or theft, driver identification for security, or rerouting in the case of traffic or inclement weather.

System Availability

Two manufacturers have announced products based on Intel® Architecture that will be available for the aftermarket later this year.

Visteon Corporation, a division of Ford Motor Company, announced at the Consumer Electronics Show (CES) in January that it will begin offering ICES, a dashboard Car PC that features hands-free cell phone control, MP3 playback, wireless Internet data, and navigation.

Clarion Corporation also announced at CES that it would offer a new version of its AutoPC based on Intel Architecture. The system offers DVD playback, wireless traffic data, navigation, and hands-free cell phone control.

Building Blocks

Today Intel offers four components that have been qualified to the extended temperature range of –45 to 85°C required for automotive applications:

- 166-MHz Pentium® processor with MMX™ Technology 166 MHz
- 430TX PCIset consisting of the 82439TX System Controller and the 82371EB PCI-to-ISA/IDE Xcelerator
- B69000ET Graphics Accelerator

Datasheets are located in the Car PC area of the Intel Developer Web site.

Summary

PC applications to keep drivers safe, entertained, and productive are already being integrated into automobiles. Products coming to market as factory equipment and as aftermarket options will assist with navigation, voice and data communication, and more. Several Intel® components have already been qualified for use in these advanced functions, and more are under evaluation.

More Info

For more information on Intel's developments in the Car PC market, including datasheets on products qualified for automotive applications, visit Intel's Car PC Web site.

For information on Visteon products, visit the Visteon Web site.

For information on Clarion products, visit the company's AutoPC Web site.

Author Bio

Lee Machen is a technical marketing engineer in the Handheld Components Division in Arizona, and is responsible for North American customer support and development. He has been with Intel since 1996. In previous positions Lee worked in tape-BGA package development and qualification and in die sales for the Embedded Microcomputer Division. Before joining Intel, Lee was a product engineer for Symbios Logic in Colorado. He holds a B.S.E.E. from the University of Michigan and will complete his M.B.A. from Arizona State University in May.

—End of Intel Developer Update Magazine Issue 7—